

WHAT IS CLAIMED IS:

1. A method for controlling a flow of packet data in a memory management unit of a network switch fabric, said method comprising:

receiving a first portion of a data packet at a port on an ingress bus ring of the network switch fabric;

5 determining a class of service for the data packet based on said first portion;

storing said portion in a packer RAM of the port based on said class of service;

storing subsequent portions of the data packet in the packer RAM;

10 sending a predetermined number of portions to a packet pool RAM, once the predetermined number of portions have been received;

10 sending a reference pointer to a first predetermined number of portions to a transaction queue once an end of packet is detected;

detecting a presence of a ready packet in the transaction queue by an egress scheduler and notifying an unpacker of the ready packet;

putting the ready packet into a FIFO through the unpacker; and

15 sending the ready packet to an ingress/egress module.

2. A method as recited in claim 1, wherein the step of sending a predetermined number of portions comprises the steps of:

waiting for the predetermined number of portions to be received;

accessing a pointer to a current free portion of the packet pool RAM by a link list

5 array;

writing the predetermined number of portions to the packet pool RAM by a link list array using the pointer;

setting or updating a link list for the data packet and the pointer; and
repeating the above process steps for subsequent portions of the data packet.

3. A method as recited in claim 1, wherein said step of notifying the unpacker of the ready packet further comprises determining a next packet to be the ready packet from a plurality of packets received by a plurality of ports and stored in the packet pool memory according to predetermined priority rules.

4. A method as recited in claim 3, wherein the predetermined priority rules comprise a deficit round robin scheduling algorithm.

5. A method as recited in claim 3, wherein the predetermined priority rules comprise a weighted round robin scheduling algorithm.

6. A method as recited in claim 1, wherein said step of detecting the presence of the ready packet further comprises the step monitoring an age of packet data in the transaction queue and purging the packet data when the age is greater than a predetermined value.

7. A method as recited in claim 1, wherein the step of sending a reference pointer further comprises determining if a class of service class in the transaction queue has reached a limit and purging the data packet when the transaction queue has reached the limit.

8. A memory management unit for a network switch fabric comprising:

means for receiving a first portion of a data packet at a port on an ingress bus ring of the network switch fabric;

means for determining a class of service for the data packet based on said first portion;

means for storing said portion in a packet RAM of the port based on said class of service;

means for sending a predetermined number of portions to a packet pool RAM;

means for sending a reference pointer to a first predetermined number of portions to a transaction queue;

means for detecting a presence of a ready packet in the transaction queue by an egress scheduler and means for notifying an unpacker of the ready packet;

means for putting the ready packet into a FIFO through the unpacker; and

means for sending the ready packet to an ingress/egress module;

wherein the means for sending a predetermined number of portions sends portions to the packet pool RAM once the predetermined number of portions have been received and the means for sending a reference pointer send the reference pointer once an end of packet is detected.

9. A memory management unit as recited in claim 8, wherein the means for sending a predetermined number of portions comprises:

means for determining whether the predetermined number of portions have been received;

means for accessing a pointer to a current free portion of the packet pool RAM by

a link list array;

means for writing the predetermined number of portions to the packet pool RAM
by a link list array using the pointer; and

means for setting or updating a link list for the data packet and the pointer.

10. A memory management unit as recited in claim 8, wherein said means for notifying the unpacker of the ready packet further comprises means for determining a next packet to be the ready packet from a plurality of packets received by a plurality of ports and stored in the packet pool memory, according to predetermined priority rules.

11. A memory management unit as recited in claim 10, wherein the predetermined priority rules used by said means for determining comprise a deficit round robin scheduling algorithm.

12. A memory management unit as recited in claim 10, wherein the predetermined priority rules used by said means for determining comprise a weighted round robin scheduling algorithm.

13. A memory management unit as recited in claim 8, wherein said means for detecting the presence of the ready packet further comprises means for monitoring an age of packet data in the transaction queue and means for purging the packet data when the age is greater than a predetermined value.

14. A memory management unit as recited in claim 8, wherein the means for

17. A method as recited in claim 15, wherein said step of determining an egress port bitmap comprises the steps of:

determining if a destination for the incoming data packet is a trunked port;

reading from a trunk group table, and

5 deriving the egress port bitmap based on entries in the trunk group table.

18. A method as recited in claim 15, wherein said particular packet information includes an opcode value that identifies whether the incoming data packet is a unicast packet, a multicast packet, a broadcast packet or resulted in a destination lookup failure.

19. A network switch fabric for forwarding data comprising:

means for receiving an incoming data packet at a first port of the fabric;

means for reading a first packet portion, less than a full packet length, to determine particular packet information, said particular packet information including a source address and a destination address;

means for determining an egress port bitmap based on a lookup in a forwarding table;

means for determining if the destination address belongs to a trunk group of trunked ports;

means for determining a particular trunked port of the trunk group; and

means for forwarding the incoming data packet based on the egress port bitmap or the particular trunked port,

wherein the means for forwarding the incoming data packet forwards the incoming data packet based on the egress port bitmap when the destination address

15 does not belong to the trunk group, and forwards the incoming data packet based on the particular trunked port when the destination address does belong to the trunk group.

20. A network switch fabric as recited in claim 19, wherein said means for determining the particular trunked port of the trunk group comprises means for calculating a hash value based on the source address and the destination value and means for selecting the particular trunked port based on the hash value.

21. A network switch fabric as recited in claim 19, wherein said particular packet information includes an opcode value that identifies whether the incoming data packet is a unicast packet, a multicast packet, a broadcast packet or resulted in a destination lookup failure.

22. A network switch fabric as recited in claim 19, further comprising means for determining a class of service for the incoming data packet from said particular packet information and means for setting a priority for forwarding based on the class of service.

23. A network switch fabric for forwarding data comprising:

means for receiving an incoming data packet at a first port of the fabric;

means for reading a first packet portion, less than a full packet length, to determine particular packet information, said particular packet information including a source address and a destination address;

means for determining at least one egress port based on a lookup in a forwarding table;

means for determining a class of service for the incoming data packet based on the particular packet information;

10 means for assigning data from the incoming data packet to a queue of a series of queues based on the at least one egress port and the class of service; and

means for forwarding data packets sequentially from each queue of said series of queues.

24. A network switch fabric as recited in claim 23, wherein the means for assigning data from the incoming data packet to a queue comprises means for packing the data from the incoming data packet into cells having a specific cell-length and means for assigning the packed data to the queue.

25. A network switch fabric as recited in claim 23, wherein said means for determining an egress port bitmap comprises:

means for determining if a destination for the incoming data packet is a trunked port;

5 means for reading from a trunk group table, and

means for deriving the egress port bitmap based on entries in the trunk group table.

26. A network switch fabric as recited in claim 23, wherein said particular packet information includes an opcode value that identifies whether the incoming data packet is a unicast packet, a multicast packet, a broadcast packet or resulted in a destination lookup failure.